DOCUMENT RESUME

ED 414 934 IR 056 792

AUTHOR	Kenney, Anne R.
TITLE	Digital Image Quality: From Conversion to Presentation and Beyond.
PUB DATE	1997-04-00
NOTE	16p.; Paper presented at the Conference on Scholarly Communication and Technology (Atlanta, GA, April 24-25, 1997), see IR 056 774.
AVAILABLE FROM	Association of Research Libraries (ARL) Web site: http://www.arl.org/scomm/scat/
PUB TYPE	Reports - Descriptive (141) -- Speeches/Meeting Papers (150)
EDRS PRICE	MF01/PC01 Plus Postage.
DESCRIPTORS	Access to Information; *Benchmarking; Cost Effectiveness; Design Requirements; *Evaluation Methods; Information Dissemination; Information Storage; Information Technology; *Quality Control
IDENTIFIERS	*Digital Imagery; Digital Technology

ABSTRACT

This paper advocates a strategy of "full informational capture" to ensure that digital objects rich enough to be useful over time are created in the most cost-effective manner. Digital benchmarking is a systematic procedure to forecast a likely outcome. The benchmarking approach can be applied across the full continuum of the digitization chain, from conversion to storage to access to presentation. It offers a means to evaluate choices for how best to balance quality, costs, timeliness, user requirements, and technological capabilities in the conversion, delivery, and maintenance of digital resources. This paper discusses how benchmarking works, including objective and subjective evaluation; determining scanning resolution requirements for replacement purposes; conversion benchmarking beyond text; benchmarking display requirements; the physical limitations of computer monitors; and how benchmarking for display works. (AEF)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED 414 934

Scholarly Communication and Technology

Conference Organized by The Andrew W. Mellon Foundation

at Emory University
April 24-25, 1997

Copyright © of the papers on this site are held by the individual authors or The Andrew W. Mellon Foundation. Permission is granted to reproduce and distribute copies of these works for nonprofit educational or library purposes, provided that the author, source, and copyright notice are included on each copy. For commercial use, please contact **Richard Ekman** at the The Andrew W. Mellon Foundation.

Session #5 Technical Choices and Standards

Digital Image Quality: From Conversion to Presentation and Beyond

Anne R. Kenney
Associate Director, Department of Preservation
Cornell University Library

PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
Richard Ekman
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
☐ This document has been reproduced as received from the person or organization originating it.
☐ Minor changes have been made to improve reproduction quality.
• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

As this conference attests, there are a number of significant digital library projects underway designed to test the economic value of digital over physical library building. Business cases are being developed to demonstrate the economics of digital applications to assist research and cultural institutions respond to the challenges of the information explosion, spiraling storage and subscription costs, and increasing user demands. These projects also reveal that the costs of selecting, converting, and making digital information available can be staggering, and that the costs of archiving and migrating that information over time are not insignificant.

Economic models comparing the digital to the traditional library show that digital will become more cost-effective *provided* the following four assumptions prove true:

- that digital collections can alleviate the need to support full traditional libraries at the local

IR056793

BEST COPY AVAILABLE

level,

- that use will increase with electronic access, and
- that the long-term value of digital collections will exceed the costs associated with their creation, maintenance, and delivery.[1]

These four assumptions--resource sharing, lower costs, meeting user demands for timely and enhanced access, and continuing value of information--presume that electronic files will have relevant content and meet baseline measures of functionality over time. Although a number of conferences and publications have addressed the need to develop selection criteria for digital conversion, and to evaluate the effective use of digitized material, more rhetoric than substantive information has emerged regarding the impact on scholarly research of creating digital collections and making them accessible over networks.

As has been argued elsewhere, I believe that digital conversion efforts will prove economically viable only if they focus on creating electronic resources for long-term use. Retrospective sources should be selected carefully based on their intellectual content; digital surrogates should effectively capture that intellectual content; and access should be more timely, usable, or cost-effective than is possible with original source documents. In sum, I would argue that long-term utility should be defined by the informational value and functionality of digital images, not limited by technical decisions made at the point of conversion or anywhere else along the digitization chain. In this paper, I advocate a strategy of "full informational capture" to ensure that digital objects rich enough to be useful over time are created in the most cost-effective manner.[2]

There is much to be said for capturing the best possible digital image you can. From a preservation perspective, the advantages are obvious. An "archival" digital master can be created to replace rapidly deteriorating originals or to reduce storage costs and increase access times to office back files, provided the digital surrogate is a trusted representation of the hardcopy source. It also makes economic sense, as Michael Lesk has noted, to "turn the pages once" and produce a sufficiently high level image so as to avoid the expense of reconverting at a later date when technological advances require, or can effectively utilize, a richer digital file.[3] This economic justification is particularly compelling as the labor costs associated with identifying, preparing, inspecting, and indexing digital information far exceed the costs of the scan itself. In recent years, the costs of scanning and storage have declined rapidly, narrowing the gap between high quality and low quality digital image capture. Once created, the archival master can then be used to create derivatives to meet a variety of current and future user's needs: high resolution may be required for printed facsimiles, on-screen detailed study, [4] and in the future for intensive image processing; moderate to high resolution for character recognition systems and image summarization techniques;[5] and lower resolution images, encoded text, or PDFs derived from the digital masters for on-screen display and browsing.[6] The quality, utility, and expense of all these derivatives will be directly affected by the quality of the initial scan.[7]

If there are compelling reasons for creating the best possible image, there is also much to be said for not capturing more than you need. At some point, adding more resolution will not result in greater quality, just a larger file size and higher costs. The key is to match the conversion process to the informational content of the original. At Cornell, we've been investigating digital imaging in a preservation context for eight years. For the first three years, we concentrated on

what was technologically possible--on determining the best image capture we could secure. For the last five years, we've been striving to define the minimal requirements for satisfying informational capture needs. No more, no less.

Digital Benchmarking

To help us determine what is minimally acceptable, we have been developing a methodology, called benchmarking. Digital benchmarking is a systematic procedure to forecast a likely outcome. It begins with a assessment of the source documents and user needs; factors in relevant objective and subjective variables associated with stated quality, cost, and/or performance objectives; involves the use of formulas that represent the inter-relationship of those variables to desired outcomes; and concludes with confirmation through carefully structured testing and evaluation. If the benchmarking formula does not consistently predict the outcome, it may not contain the relevant variables or reflect their proper relationship--and it should be revised.

Benchmarking does not provide easy answers, but a means against which to evaluate possible answers for how best to balance quality, costs, timeliness, user requirements, and technological capabilities in the conversion, delivery, and maintenance of digital resources. It is also intended as a means to formulate a range of possible solutions on the macro level rather than on an individual, case-by-case basis. For many aspects of digital imaging, benchmarking is still unchartered territory. Much work remains to be able to define conversion requirements for certain document types, e.g., photographs and high end book illustrations; for conveying color information; for evaluating the effects of new compression algorithms; and for providing access on a mass scale to a digital database of material representing a wide range of document types and document characteristics.

We began benchmarking with the conversion of printed text. We anticipate that within 2 years, quality benchmarks for image capture and presentation of the broad range of paper and film based research materials--including manuscripts, graphic art, halftones, and photographs--will be well defined through a number of projects currently underway.[8] In general, these projects are designed to be system independent and are based increasingly on assessing the attributes and functionality characteristic of the source documents themselves, coupled with an understanding of user perceptions and requirements.

Why benchmarking?

Because there are no standards for image quality, because different document types require different scanning processes, there is no "silver bullet" for conversion. This frustrates many librarians and archivists who are seeking a simple solution to a complex issue. I suppose if there really were the need for a silver bullet, I'd recommend that most source documents be scanned at a minimum of 600 dpi with 24 bit color, but that would result in tremendously large file sizes, and a hefty conversion cost. One would also be left with the problems of transmitting and displaying those images.

We began benchmarking with conversion, but we are now applying this approach to the presentation of information on screen. The number of variables that govern display are many, and it will come as no surprise that they preclude the establishment of a single best method for

presenting digital images. But here, too, the urge is strong to seek a single solution. If display requirements paralleled conversion requirements--that is, if a 600 dpi, 24 bit image had to be presented on screen, then *at best*, with the highest resolution monitors commerically available, only documents whose physical dimensions did not exceed 2.7" x 2.13" could be displayed--and they could not be displayed at their native size. Now most of us are interested in converting and displaying items that are larger than postage stamps, so these "simple solutions" are for most purposes impractical, and compromises will have to be made.

The object of benchmarking is to make informed decisions about a range of choices and to understand in advance the consequences of such decisions. The benchmarking approach can be applied across the full continuum of the digitization chain, from conversion to storage to access to presentation. Our belief at Cornell is that benchmarking must be approached holistically, that it is essential to understand at the point of selection what the consequences downstream for conversion and presentation will be. This is especially important as institutions consider inaugurating large scale conversion projects. Towards this end, the advantages of benchmarking are several in number.

1. Benchmarking is first and foremost a management tool, designed to lead to informed decision-making. It offers a starting point and a means for narrowing the range of choices to a manageable number. Although clearly benchmarking decisions must be judged through actual implementations, the time spent in experimentation can be reduced, the temptaton to overstate or understate requirements may be avoided, and the initial assessment requires no specialized equipment nor expenditure of funds. Benchmarking allows one to scale knowledgeably, to make decisions on a macro level, rather than to determine those requirements through item-by-item review or by setting requirements for groups of materials that may be adequate for only a portion of them.

2. Benchmarking provides a means for interpreting vendor claims. If you have spent any time reading product literature, you may have become convinced, as I have, that the sole aim of any company is to sell its product. Technical information will be presented in the most favorable light, which is often incomplete and intended to discourage product comparisons. One film scanner for instance may be advertised as having a resolution of 7500 dpi; another may claim 400 dpi. In fact, these two scanners could provide the very same capabilities but it may be difficult to reach that conclusion without additional information. You may end up spending considerable time on the phone, first getting past the marketing representatives, and then questioning closely those with a technical understanding of the product's capabilities. If you have benchmarked your requirements, you will be able to focus the discussion on your particular needs.

3. Benchmarking can assist you in negotiating with vendors for services and products. I've spent many years advocating the use of 600 dpi bitonal scanning for printed text and invariably when I begin a discussion with a representative of an imaging service bureau, he will try to talk me out of that high a resolution, claiming that I do not need it or that it will be exhorbitantly expensive. I suspect he is in part motivated to make those claims because he believes them, and in part because his company may not provide that service and he wants my business. If I had not benchmarked my resolution requirements, I might be pursuaded by what this salesperson has to say.

4. Benchmarking can lead to careful management of resources. If you know up front what your requirements are likely to be and the consequences of those requirements, you can develop a

budget that reflects the actual costs, identify prerequisites for meeting those needs, and, perhaps most important, avoid costly mistakes. Nothing will doom an imaging project more quickly than buying the wrong equipment or having to manage image files that are not supported by your institution's technical infrastructure.

5. Benchmarking can also allow you to predict what you can deliver under specific conditions. It is important to understand that an imaging project may break at the weakest link in the digitization chain. For instance, if your institution is considering scanning its map collection, one should be realistic about what ultimately can be delivered to the user at her desktop. Benchmarking lets you predict how much of the image and what level of detail contained therein can be presented on-screen for various monitors. Even with the most expensive monitor available today, presenting oversize material completely, with small detail intact, is impractical.

How Does It Work?

Having spent some time extolling the virtues of digital benchmarking, I'd like to turn next to describing this methodology as it applies to conversion, and then to move to a discussion of on-screen presentation.

Objective Evaluation:
Determining what constitutes informational content becomes the first step in the conversion benchmarking process. This can be done objectively or subjectively. Let's consider an objective approach first. One way to do this would be to peg conversion requirements to the process used to create the original document. Take resolution, for instance. Film resolution can be measured by the size of the silver grains suspended in an emulsion, whose distinct characteristics are appreciated only under microscopic examination. Should we aim for capturing the properties of the chemical process used to create the original? Or should we peg resolution requirements at the recording capability of the camera or printer used?

There are objective scientific tests that can measure the overall information carrying capacity of an imaging system, such as the Modulation Transfer Function, but such tests require expensive equipment and are still beyond the capabilities of most outside industry or research labs. In practical applications, the resolving power of a microfilm camera is measured by means of a technical test chart where the distinct number of black and white lines discerned is multiplied by the reduction ratio used to determine the number of line pairs per millimeter. A system resolution of 120 line pairs per millimeter is considered good; above 120 is considered excellent. To capture digitally all the information present on a 35mm frame of film with a resolution of 120 lppm would take a bitonal film scanner with a pixel array of 12,240.[9] There is no such beast on the market today.

How far down this path should we go? It may be appropriate to require that the digital image accurately depict the gouges of a wood cut or the scoops of a stipple engraving, but what about the exact dot pattern and screen ruling of a halftone? the strokes and acid bite of an etching? the black lace of an aquatint that only becomes visible at a magnification above 25x? Offset publications are printed at 1200 dpi--should we chose that resolution as our starting point for scanning text? Significant information may well be present at that level in some cases, as may be argued for medical x-rays, but in other cases, attempting to capture all possible information will far exceed the inherent properties of the image as distinct from the medium and process used to create it. Consider for instance a 4 *x* 5 negative of a badly blurred photograph. The negative is

incredibly information dense, but the information it conveys is not significant.

Obviously, any practical application of digital conversion would be overwhelmed by the recording, computing, and storage requirements that would be needed to support capture at the structure or process level. Although offset printing may be produced at 1200 dpi, most individuals would not be able to discern the difference between a 600 dpi and a 1,000 dpi digital image of that page, even under magnification. In choosing the higher resolution one would be adding more bits, increasing the file size, but with little to no appreciable gain. The difference between 300 dpi and 600 dpi, however, can be easily observed, and, in my opinion, is worth the extra time and expense to obtain. The relationship between resolution and image quality is not linear: at some point as resolution increases, the gain in image quality will level off. Benchmarking will help you to determine where the leveling begins.

Subjective Evaluation:
I would argue, then, that determining what constitutes informational content is best done subjectively. It should be based on an assessment of the attributes of the document rather than the process used to create that document. Reformatting via digital--or analog--techniques presumes that the essential meaning of an original can somehow be captured and presented in another format. There is always some loss of information when an object is copied. The key is to determine whether that informational loss is significant or not. Obviously for some items, particularly those of intrinsic value, a copy can only serve as a surrogate, not as a replacement. This determination should be made by those with curatorial responsibility and a good understanding of the nature and signficance of the material. Those with a trained eye should consider the attributes of the document itself as well as the immediate and potential uses that researchers will make of its informational content.

Determining Scanning Resolution Requirements For Replacement Purposes:

To illustrate benchmarking for conversion, let's consider the brittle book. For brittle books published during the last century and a half, detail has come to represent the size of the smallest significant character in the text, usually the lower case "e." To capture this information--which consists of black ink on a light background--resolution is the key determinant of image quality.

Benchmarking resolution requirements in a digital world has its roots in micrographics, where standards for predicting image quality are based on the Quality Index (QI). QI provides a means for relating system resolution and text legibility. It is based on multiplying the height of the smallest significant character "h" by the smallest line pair pattern resolved by a camera on a technical test target, "p," QI=h x p. The resulting number is called the Quality Index, and it is used to forecast levels of image quality--marginal (3.6), medium (5.0) or high (8.0)--that will be achieved on the film. This approach can be used in the digital world, but a number of adjustments must be made to account for the differences in the ways in which microfilm cameras and scanners capture detail.[10] Specifically, it is necessary to:

1. Establish levels of image quality for digitally rendered characters that are analogous to those established for microfilming (illustration showing differences in quality degradation). Note that in photographically reproduced images, quality degradation results in a fuzzy or blurred image. Usually degradation with digital conversion is revealed in the ragged or stairstepped appearance of diagonal lines or curves, known as aliasing or "jaggies."

2. Rationalize system measurements. Digital resolution is measured in dots per inch; classic resolution is measured in line pairs per millimeter. To calculate QI based on scanning resolution, one must convert from one to the other. One millimeter equals .039 inches, so to determine the number of dots per millimeter, you will need to multiply the DPI by .039.

3. Equate dots to line pairs. Again, classic resolution refers to line *pairs* per millimeter (one black line and one white line), and since a dot occupies the same space as a line, two dots must be used to represent one line pair. This means the dpi must be divided by two to be made equivalent to "p."

With these adjustments, we can modify the QI formula to create a digital equivalent. From QI= px h,

$$\text{we now have } QI = \frac{.039dpi}{2} \times h$$

which can be simplified to .0195dpi x h.

For bitonal scanning, we would also want to adjust for possible misregistration due to sampling errors brought about in the thresholding process in which all pixels are reduced to either black or white. To be on the conservative side, the authors of AIIM TR26-1993 advise increasing the input scanning resolution by at least 50% to compensate for possible image detector mis-alignment. The formula would then be

$$QI = \frac{.039dpi}{3} \times h \quad \text{which can be simplified to } .013dpi \times h.$$

So how does all this work?

Consider a printed page that contains characters measuring 2mm high and above. If the page were scanned at 300 dpi, what level of quality would you expect to obtain? By plugging in the dpi and the character height and solving for QI, you would discover that you can expect a QI of 8, or excellent rendering.

One can also solve the equation for the other variables. Consider for example a scanner with a maximum of 400 dpi. You can benchmark the size of the smallest character that you could capture with medium quality (a QI of 5), which would be .96mm high. Or you can calculate the input scanning resolution required to achieve excellent rendering of a character that is 3 mm high (200 dpi).

With this formula, and an understanding of the nature of your source documents, you can benchmark the scanning resolution needs for printed material. We took this knowledge and applied it to the types of documents we were scanning--brittle books published from 1850-1950. We reviewed printers' type sizes commonly used by publishers during this period, and discovered that virtually none utilized type fonts smaller than 1 mm in height, which, according to our benchmarking formula, could be captured with excellent quality using 600 dpi bitonal scanning. We then tested these benchmarks by conducting an extensive on-screen and in print examination of digital facsimiles for the smallest font-sized Roman and non-Roman type scripts used during this period. This verification process confirmed that an input scanning resolution of 600 dpi was indeed sufficient to capture the monochrome text-based information contained in

virtually all books published during the period of paper's greatest brittleness. Although many of those books do not contain text that is as small as 1 mm in height, a sufficient number of them do. To avoid the labor and expense of performing item by item review, we currently scan all books at 600 dpi resolution.[111]

Conversion Benchmarking Beyond Text

Although we've conducted most of our experiments on printed text, we are beginning to benchmark resolution requirements for non-textual documents as well. For non-text based material, we have begun to develop a benchmarking formula that would be based on the width of the smallest stroke or mark on the page rather than a complete detail. This approach was used by the Nordic Digital Research Institute to determine resolution requirements for the conversion of historic Icelandic maps, and is being followed in the current New York State Kodak Photo CD project being conducted at Cornell on behalf of the Eleven Comprehensive Research Libraries of New York State. The measurement of such fine detail will require the use of a 25-50x lupe with a metric hairline that differentiates below .1mm.

Benchmarking for conversion can be extended beyond resolution to tonal reproduction (both grayscale and color), to the capture of depth, overlay, and translucency, to assessing the effects of compression techniques and levels of compression used on image quality, to evaluating the capabilities of a particular scanning methodology, such as the Kodak Photo CD format. It can also be used for evaluating quality requirements for a particular category of materials, e.g., halftones, or to examine the relationship between the size of the document and the size of its significant details, a very challenging relationship which affects both the conversion and the presentation of maps, newspapers, architectural drawings, and other oversized, highly detailed source documents.

Benchmarking involves both subjective and objective components. There must be the means to establish levels of quality (through technical targets, samples of acceptable materials), the means to identify and measure significant information present in the document, the means to relate one to another via a formula, and the means to judge results on-screen and in print for a sample group of documents. Armed with this information, benchmarking enables informed decision making--which often leads to a balancing act involving tradeoffs between quality and cost, between quality and completeness, between completeness and size, or quality and speed.

Benchmarking Display Requirements:

Quality assessments can be extended beyond capture requirements to the presentation and timeliness of delivery options. We begin our benchmarking for conversion with the attributes of the *source documents*. We begin our benchmarking for display with the attributes of the *digital images*.

I believe that all researchers in their heart of hearts expect three things from displayed digital images: they want the full size image to be presented on screen; they expect legibility and adequate color rendering, and they want images to be displayed quickly. Of course they want lots of other things, too, such as the means to manipulate, annotate, and compare images, and for text-based material, they want to be able to conduct key word searches across the images. But for the moment, let's just consider those three requirements: full image, full detail and tonal

reproduction, quick display.

Unfortunately, for many categories of documents, satisfying all three criteria at once will be a problem, given the limitations of screen design, computing capabilities, and network speeds. Benchmarking screen display must take all these variables into consideration *and* the attributes of the digital images themselves as user expectations are weighed one against the other. We are just beginning to investigate this interrelationship at Cornell, and although our findings are still tentative and not broadly confirmed through experimentation, I'm convinced that display benchmarking will offer the same advantages as conversion benchmarking to research institutions that are beginning to make their materials available electronically.[12]

Now for the good news: it is easy to display the complete image and it is possible to display it quickly. It is easy to ensure screen legibility--in fact intensive scrutiny of highly detailed information is facilitated on screen. Color fidelity is a little more difficult to deliver, but progress is occurring on that front.[13]

Now for the not so good news: given common desktop computer configurations, it may not be possible to deliver full 24-bit color to the screen--the monitor may have the native capability but not enough video memory or its refresh rate can not sustain a non-flickering image. The complete image that is quickly displayed may not be legible. A highly detailed image may take a long time to deliver and only a small percent of it will be seen at any given time. You may call up a photograph of Yul Brenner only to discover you have landed somewhere on his bald pate.

Benchmarking will allow you to predict in advance the pros and cons of digital image display. Conflicts between legibility and completeness, between timeliness and detail, can be identified and compromises developed. Benchmarking allows you to predetermine a set process for delivering images of uniform size and content, and to assess how well that process will accommodate other document types. Scaling to 72 dpi and adding 3 bits of gray may be a good choice for technical reports produced at 10 point type and above, but will be totally inadequate for delivering digital renderings of full-size newspapers.

To illustrate benchmarking as it applies to display, consider the first two user expectations: complete display and legibility. We expect printed facsimiles produced from digital images to look very similar to the original. They should be the same size, preserve the layout, and convey detail and tonal information that is faithful to the original. Many readers assume that the digital image on screen can also be the same, that if the page were correctly converted, it could be brought up at approximately the same size and with the same level of detail as the original. It is certainly possible to scale the image to be the same size as the original document, but chances are information contained therein will not be legible.

If the scanned image's dpi does not equal the screen dpi, then the image on-screen will either appear larger or smaller than the original document's size. Because scanning dpi most often exceeds the screen dpi, the image will appear larger on the screen--and chances are not all of it will be represented at once. This is because monitors have a limited number of pixels that can be displayed both horizontally and vertically. If the number of pixels in the image exceed those of the screen and the scanning dpi is higher, the image will be enlarged on the screen and not completely presented.

The problems of presenting completeness, detail, and native size are more pronounced in display than in printing. In the latter, industry is capable of very high printing resolutions, and the total

number of dots that can be laid down for a given image is great, enabling the creation of facsimiles that are the same size--and often with the same detail--as the original.

The limited pixel dimensions and dpi of monitors can be both a strength and a weakness. On the plus side, detail can be presented more legibly and without the aid of a microscope, which for those conducting extensive textual analysis may represent a major improvement over reviewing the source documents themselves. For instance, papyrologists can rely on monitors to provide the enlarged view of fragment details required in their study. When the originals themselves are examined, they are typically viewed under a microscope at 4 to 10x magnification.[14] Art historians can zoom in on high resolution images to enlarge details or to examine brush strokes that convey different surfaces and materials.[15] On the down side, because the screen dpi is often exceeded by the scanning dpi, and screens have very limited pixel dimensions, many documents can not be fully displayed *if* legibility must be conveyed. This conflict between overall size and level of detail is most apparent when dealing with oversized material, but it also affects a surprisingly large percentage of normal-sized documents as well.

Consider the physical limitations of computer monitors:

Typical monitors offer resolutions from 640 x 480 at the low end to 1600 x 1200 at the high end. The lowest level SVGA monitor offers the possibility of displaying material at 1024 x 768. These numbers, known as the pixel matrix, refer to the number of horizontal by vertical pixels painted on the screen when an image appears.

In product literature, monitor resolutions are often given in dpi which can range from 60 to 120, depending on the screen width and horizontal pixel dimension. The screen dpi can be a misleading representation of a monitor's quality and performance. For example, when SVGA resolution is used on a 14", 17", and 21" monitor, the screen dpi decreases as screen size increases. We might intuitively expect image resolution to increase with the size of the monitor, not decrease. In reality the same amount of an image--and level of detail--would be displayed on all three monitors set to the same pixel dimensions. The only difference would be that the image displayed on the 21 inch monitor would appear enlarged compared to the same image displayed on the 17 and 14 inch monitors.

The pixel matrix of a monitor limits the number of pixels of a digital image that can be displayed at any one time. And, if there is insufficient video memory, you will also be limited to how much gray or color information can be supported at any pixel dimension. For instance, while the three-year old 14" SVGA monitor on my desk supports a 1024 x 768 display resolution, it came bundled with half a megabyte of video memory. It can not display an 8-bit grayscale image at that resolution and it can not display a 24 bit color image at all, even if it is set at the lowest resolution of 640 x 480. Even if I increased its VRAM, I would be bothered by an annoying flicker, as the monitor's refresh rate is not great enough to support a stable image on screen at higher resolutions. It is not coincidental that while the most basic SVGA monitors can support a pixel matrix of 1024 x 768, most of them come packaged with the monitor set at a resolution of 800 x 600. As others have noted, network speeds and the limitations of graphical user interfaces will also affect profoundly user satisfaction with on-screen presentation of digital images.

So how does benchmarking for display work?

Consider the brittle book and how best to display it. Recall that it may contain font sizes at 1 mm and above, so we have scanned each page at 600 d pi, bitonal mode. Let's assume that the typical page averages 4" x 6" in size. The pixel matrix of this image will be: 4 x 600 by 6 x 600, or 2400 x 3600--far above any monitor pixel matrix currently available. Now if I want to display that image at its full scanning resolution on my monitor, set to the default resolution of 800 x 600, it should be obvious to many of you that I will be showing only a small portion of that image--approximately 5% of it will appear on the screen. Let's suppose I went out and purchased a $2,500 monitor that offered a resolution of 1600 x 1200. I'd still only be able to display less than a fourth of that image at any one time.

Obviously for most access purposes, this display would be unacceptable. It requires too much scrolling or zooming out to study the image. If it is an absolute requirement that the full image be displayed with all details fully rendered, I'd suggest converting only items whose smallest significant detail represents nothing smaller than one third of 1% of the total document surface. This means that if you had a document with a one millimeter high character that was scanned at 600 dpi and you wanted to display the full document at its scanning resolution on a 1024 x 768 monitor, the document's physical dimensions could not exceed 1.7" (horizontal) x 1.3" (vertical). This may work well for items such as papyri which are relatively small, at least as they have survived to the present. It also works well for items that are physically large *and* contain large-sized features, such as posters that are meant to be viewed from a distance. If the smallest detail on the poster measured one inch, the poster could be as large as 42" x 32" and still be fully displayed with all detail intact.[16]

Most images will have to be scaled down from their scanning resolutions for on screen access, and this can occur a number of ways. Let's first consider full display on the monitor, and then consider legibility. In order to display the full image on a given monitor, the image pixel matrix must be reduced to fit within the monitor's pixel dimensions. The image is scaled by setting one of its pixel matrixes to the corresponding pixel dimension of the monitor.[17]

To fit the complete page image from our brittle book on a monitor set at 800 *x* 600, we would scale the vertical dimension of our image to 600; the horizontal dimension would be 400 to preserve the aspect ratio of the original. By reducing the 2400 x 3600 pixel image to 400 x 600, we will have discarded 97% of the information in the original. The advantages to doing this are several: it facilitates browsing by displaying the full image, it decreases file size which in turn decreases the transmission time. The down side should also be obvious. There will be a major decrease in image quality as a significant number of pixels are discarded. In other words, the image can be fully displayed, but the information contained in that image may not be legible. To determine whether that information will be useful, we can turn to the use of benchmarking formulas for legible display:

Benchmarking resolution formulas for scaling bitonal and grayscale images for on-screen display:[18]

dpi = QI/(.03h)

QI = dpi *x*.03h

h = QI/(.03dpi)

Note: Recall that in the benchmarking resolution formulas for conversion, dpi refers to the

scanning resolution. In the scaling formulas, dpi refers to the image dpi (not to be confused with the monitor's dpi).

Let's return to the example of our 4 x 6 brittle page.

If we assume we need to be able to read the 1 mm high character, but that it doesn't have to be fully rendered, then we set our QI requirement at 3.6, which should ensure legibility of characters in context. We can use the benchmarking formula to predict the scaled image dpi:

dpi = QI/.03h, or

dpi = 3.6/(.03x 1), or

dpi =120

The image could be fully displayed with minimal legibility on a 120 dpi monitor. The pixel dimensions for the scaled image would be 120 *x* 4 by 120 *x* 6, or 480 *x* 720. This full image could be viewed on SVGA monitors set at 1024 x 768 or above; slightly over 80% of it could be viewed on my monitor set at 800 x 600.

We can also use this formula to determine a preset scaling dpi for a group of documents to be conveyed to a particular clientele. Consider a scenario where your primary users have access to monitors that can support effectively an 800 x 600 resolution. We could decide whether the user population would be satisfied with receiving only 80% of the document if it meant that they could read the smallest type, which may occur only in footnotes. If your users are more interested in quick browsing, you might want to benchmark against the body of the text, rather than the smallest typed character. For instance, if the main text were in 12 point type and the smallest "e" measured 1.6 mm in height, then our sample page could be sent to the screen with a QI of 3.6 at a pixel dimension of 300 *x* 450, or an image dpi of 75--well within the capabilities of the 800 x 600 monitor.

One can also benchmark the time it will take to deliver this image to the screen--if your clientele are connected via ethernet, this image (with 3 bits of gray added to smooth out rough edges of characters and improve legibility) could be sent to the desktop in under a second--providing readers with full display of the document, legibility of the main text, and a timely delivery. If your readers are connected to the ethernet via a 9600 baud modem, however, the image will take 42 seconds to be delivered. If the footnotes must be readable, the full text can not be delivered at once and the time it will take to retrieve the image will increase. Benchmarking allows you to identify these variables and consider the tradeoffs/compromises associated with optimizing any one of them.

Conclusion:

Benchmarking is an approach, not a prescription. It offers a means to evaluate choices for how best to balance quality, costs, timeliness, user requirements, and technological capabilities in the conversion, delivery, and presentation of digital resources. The value of this approach will best be determined by extensive field testing. We at Cornell are committed to further refinement of the benchmarking methodology, and urge others to consider its utility *before* they commit considerable resources to bringing about the brave new world of digitized information.

FOOTNOTES:

[1] Stephen Chapman and Anne R. Kenney, "Digital Conversion of Library Research Materials: A Case for Full Informational Capture," *D-Lib Magazine*, October 1996.

[2] Currently, scanning is the most cost-effective means to create digital files, and digital imaging is the only electronic format that can accurately render the information, page layout, and presentation of source documents, including text, graphics, and evidence of age and use. By producing digital images, one can create an authentic representation of the original at minimal cost, and then derive the most useful version and format (e.g., marked-up text) for transmission and use.

[3] Michael Lesk, *Image Formats for Preservation and Access. A Report of the Technology Assessment Advisory Committee to the Commission on Preservation and Access*, July 1990; see also Lesk, *Substituting Images for Books: The Economics for Libraries*, January 1996.

[4] See, Charles S. Rhyne, *Computer Images for Research, Teaching, and Publication in Art History and Related Disciplines*, Commission on Preservation and Access, January 1996, p. 4, where he argues that "with each jump in [on-screen image] quality, new uses become possible."

[5] Interesting work is being conducted at Xerox PARC on image summarization, see Francine R. Chen and Dan S. Bloomberg, "Extraction of Thematically Relevant Text from Images," to appear in SDAIR,'96, pp. 163-178.

[6] An interesting conclusion from a project on the use of art and architectural images at Cornell focused on image size guidelines to support a range of user activities. For browsing, the project staff found that images must be large enough for the user to identify the image, but small enough to allow numerous images to be viewed simultaneously--the physical size on the screen preferred by users was 1.25 to 2.25 inches square. For view images in their entirety, images were sized to fit within a 5.5 inch square; for studying, detailed views covering the entire screen were necessary, and for "authoring" presentations or other multimedia projects, users preferred images that fit in a half inch square. See Noni Korf Vidal, Thomas Hickerson, and Geri Gay, "Developing Multimedia Collection and Access Tools, Appendix V. Guidelines for the Display of Images." pp. 14-17. April 1996.

[7] A number of leading experts advocate this approach, including Michael Ester of Luna Imaging, Inc. See for example: Ester, Michael, "Digital Images in the Context of Visual Collectons and Scholarshp," *Visual Resources*, Vol X, 1990, pp. 11-24 and "Specifics of Imaging Practice," *Archives & Museum Informatics*, 1995, pp. 147-158.

[8] Roger S. Bagnall, *Digital Imaging of Papyri: A Report to the Commission on Preservation and Access*, Commission on Preservation and Access, September 1995; Janet Gertz, *Oversize*

Color Images Project, 1994-1995 Final Report of Phase I, Commission on Preservation and Access, August 1995; Picture Elements, Inc., *Guidelines for Electronic Preservation of Visual materials, Part I*, (2 March 1995), and Reilly. Michael Ester argues that an "archival image" of a photograph can not be benchmarked through calculations, but should be pegged to the "functional range of an institution's reproduction sources" see p. 11 in Ester, *Digital Image Collections: Issues and Practice*, Dec. 1996(CPA). For a critique of this approach, see Stephen Chapman and Anne R. Kenney, "Digital Conversion of Library Research Materials, A Case for Full Informational Capture," *D-Lib Magazine*, October 1996.

[2] Anne R. Kenney and Stephen Chapman, "Film Scanning," (Chapter Seven) in *Digital Imaging for Libraries and Archives*, June 1996, p. 169.

[10] ANSI/AIIM MS23-1991, *Practicefor Operational Procedures/Inspection and Quality Control of First-generaton, Silver Microfilm and Documents*, Association for Information and Image Management; ANSI/AIIM TR26-1993, *Resolution as it Relates to Photographic and Electronic Imaging*, Association for Information and Image Management; and Kenney and Chapman, *Tutorial: Digital Resolution Requirements for Replacing Text-Based Material: Methods for Benchmarking Image Quality*, Commission on Preservation and Access, April 1995.

[11] For a description of this verification process, see: Anne R. Kenney, "Digital-to-Microfilm Conversion: An Interim Preservation Solution," *Library Resources and Technical Services* (October 1993), pp. 380-401; (January 1994), pp. 87-95.

[12] A fuller explanation of the display benchmarking process is included in Kenney and Chapman, "Chapter 2", *Digital Imaging for Libraries and Archives* (June 1996), Cornell University Library, pp. 76-86.

[13] Improvements in managing color digitally may be forthcoming from an international consortium of industry leaders working to develop an electronic pre-press industry standard. Their "International Color Consortium Profile Format" is intended to represent color consistently across devices and platforms.

[14] See Peter van Minnen, "Imaging the Duke papyri," (December 1995) http://odyssey.lib.duke.edu/papyrus/texts/imaging.html, and Roger S. Bagnall, *Digital Imaging of Papyri: A Report to the Commission on Preservation and Access*, Commission on Preservation and Access, September 1995.

[15] Rhyne, *Computer Images for Research, Teaching, and Publication in Art History and Related Disciplines*, Commission on Preservation and Access, 1996, p. 5.

[16] The formula forcalculting the maximum percentage of a digital image that can be displayed on screen is:
a. If both image dimensions ≤ the correspondingpixel dimensions (pd) of the screen, 100% of the image will be displayed
b. If both image dimensions > the corresponding pixel dimensions of the screen,

$$\%\text{displayed} = \frac{\text{horiz. screen pd} \times \text{vertical screen pd}}{\text{Image's horiz. Pd x image's vertical pd}} \times 100$$

c. If one of the image's dimensions ≤ the corresponding pixel dimension of the screen,

$$\%displayed = \frac{\text{screen's opposite pixel dimension}}{\text{Image's opposite pixel dimension.}} \times 100$$

17 The formula for scaling for complete display of image on screen is:
a. When digital image aspect ratio ≤ screen aspect ratio, set image's horizontal pixel dimension to the screen's horizontal pixel dimension
b. When digital image aspect ratio is > screen aspect ratio, set image's vertical pixel dimension to the screen's vertical pixel dimension.

18 This formula presumes that bitonal images are presented with a minimum level of gray (3 bits or greater), and that filters and optimized scaling routines are used to improved image presentation.

. . . . . .

For additional information about the conference, or The Andrew W. Mellon Foundation's scholarly communication initiatives, please contact Richard Ekman. For additional information about ARL or this web site contact Patricia Brennan, ARL Program Officer at (202) 296-2296.

Return to Office of Scholarly Communication Home Page

ARL Home ARL Scholarly Communication and Technology Home Page

© Association of Research Libraries, Washington, DC
Web Design by Angelo F. Cruz
Maintained by ARL Web Administrator
Last Modified: August 1, 1997